arXiv:1603.07684v1 [cs.SY] 24 Mar 2016

MULTI-OBJECT TRACKING WITH MULTIPLE BIRTH, DEATH, AND SPAWN SCENARIOS USING A RANDOMIZED HYPOTHESIS GENERATION TECHNIQUE (R-FISST)

W. Faber and S. Chakravorty
Department of Aerospace Engineering
Texas A&M University
College Station, TX

Islam I. Hussein
Applied Defense Solutions
Columbia, MD

***Abstract*—In multi-object tracking one may encounter situations were at any time step the number of possible hypotheses is too large to generate exhaustively. These situations generally occur when there are multiple ambiguous measurement returns that can be associated to many objects. This paper contains a newly developed approach that keeps the aforementioned situations computationally tractable. Utilizing a hypothesis level derivation of the Finite Set Statistics (FISST) Bayesian recursions for multi-object tracking we are able to propose a randomized method called randomized FISST (R-FISST). Like our previous methods [1], [2], this approach utilizes Markov Chain Monte Carlo (MCMC) methods to sample highly probable hypotheses, however, the newly developed (R-FISST) can account for hypotheses containing multiple births and death within the MCMC sampling. This alleviates the burden of having to exhaustively enumerate all birth and death hypotheses and makes the method more equipped to handle spawn scenarios. We test our method on Space Situational Awareness (SSA) scenarios with spawn events.**

I. Introduction

In this paper, we present a randomized approach to the multi-object tracking problem called randomized Finite Set Statistics (R-FISST). This approach allows us to approximate the full Finite Set Statistics (FISST) Bayesian recursions, account for multiple birth and death scenarios, while keeping the problem computationally tractable. We briefly show that the FISST recursions can essentially be considered as a discrete state space Bayesian filtering problem on "Hypothesis Space" with the only input from the continuous problem coming in terms of the likelihood values of the different hypotheses. It is this perspective that allows us to develop the randomized approach where we sample the highly probable hypotheses using a Markov Chain Monte Carlo (MCMC) technique. Our previous papers, [1], [2], used similar techniques to generate hypotheses, however, they required that all possible birth and death hypotheses be considered which limited us to cases with only one birth or death for computational tractability. The newly developed R-FISST can account for multiple births and death within the MCMC hypothesis generation and alleviates the computational burden of the birth and death model. This increases the R-FISST method's ability to handle situations of mass birth and death, or spawning. This new technique is tested using Space Situational Awareness (SSA) scenarios that include birth, death, and spawning.

FISST-based multi-object detection and tracking has been developed based on the mathematical theory of finite set statistics [3], [4]. The greatest challenge in implementing FISST in real-time, which is critical to any viable SSA solution, is computational burden. The first-moment approximation of FISST is known as the Probability Hypothesis Density (PHD) approach [4], [5]. The PHD has been proposed as a computationally tractable approach to applying FISST. The PHD filter essentially finds the density of the probability of an object being at a given location, and thus, can provide information about the number of objects (integral of the PHD over the region of interest) and likely location of the objects (the peaks of the PHD). The PHD can further employ a Gaussian Mixture (GM) or a particle filter approximation to reduce the computational burden (by removing the need to discretize the state space). This comes at the expense of approximating the general FISST pdf with its first-moments [5]–[8]. The PHD filter does not attempt to solve the full FISST recursions, in particular, by considering the PHD, the filter gets rid of the data association problem inherent in these problems. In other previous work, a GM approximation was applied, not to the first-moment of the FISST pdfs, but to the original full propagation and update equations derived from FISST [9], [10]. This eliminates any information loss associated with using the first-moment PHD approximation, while at the same time increasing the computational tractability of the multi-object FISST pdfs. More recently, there has been substantive research on the so-called generalized labeled multi-Bernoulli (GLMB) filters that model the birth process as a multi-Bernoulli process and tractable implementation of the filter have been proposed based on a lookahead strategy based on the cheaper PHD filter [11], [12]. Our birth and death model is also a multi-Bernoulli process, however, our implementation is at a hypothesis level: our pdf is a weighted mixture of MT-

pdfs with varying number of targets in the component MT-pdfs. The weights are precisely the hypothesis weights and an MCMC scheme is used to sample the high probability hypothesis thereby resolving the problem of hypothesis depletion, especially when the number of hypotheses become very large, for instance in the case of a spawning event that is considered here (the number of hypothesis in such cases can explode to the billions even for a moderate number of objects).

Hypothesis Oriented MHT (HOMHT) [13]–[16], and the Track Oriented MHT (TOMHT) [17] are non-FISST approaches to the multi-object tracking problem. The MHT techniques can be divided into single-scan and multi-scan methods depending on whether the method uses data from previous times to distinguish the tracks [14], [16], [18]. We show that the MHT technique and the FISST technique are essentially the same, modulo the set-theoretic representation of multi-target pdfs in FISST (which, however, does not provide any extra information). This is made possible through the hypothesis level derivation of the tracking equations that considers the full hybrid state of the problem unlike the original MHT derivation in [13]. This allows us to identify the critical structure inherent in the FISST recursions that enables us to unify the two approaches: essentially our approach results in a mixture of hypotheses, where the hypotheses themselves are MT-pdfs with possibly different number of objects. This also allows for the computationally tractable MCMC based implementation of the full FISST recursions even when the number of hypotheses grow extremely large.

The rest of the paper is organized as follows. In Section II, we give a brief introduction to the hypothesis level derivation of the FISST equations and show the relationship to MHT. In Section III, we show how to incorporate birth and death into the MCMC based randomized hypothesis selection technique that results in the newly developed RFISST algorithm. In Section IV, we show an application of the RFISST technique to multiple spawning SSA scenarios that highlight the necessity of the randomized approach.

II. A brief overview of the Hypothesis based Derivation of the FISST equations

In this section, we give a very breif overview of the hypothesis level derivations to the full FISST recursions. We will discuss only the key points of the derivation that allow us to develop the computational technique presented later in this paper. For more detailed derivations, please refer to [1], [2].

A. Framing FISST at the Hypothesis Level

We consider first the case when the number of objects is fixed. Assume that the number of objects is M, and each object state resides in $\Re^N$. Consider some time instant $t-1$, and the data available for the multi-object tracking problem till the current time $\mathcal{F}^{t-1}$. Let H_i denote the i^{th} hypothesis at time $t-1$, and let $\{X\}$ denote the underlying continuous state. For instance, given the $N-$ object hypothesis, the underlying state space would be $\{X\} = \{X_1, X_2, \cdots X_N\}$ where X_j denotes the state of the j^{th} object under hypothesis H_i and resides in $\Re^N$. Let $p(\{X\}, i/\mathcal{F}^{t-1})$ denote the joint distribution of the state-hypothesis pair after time $t-1$. Using the rule of conditional probability:

$$p(\{X\}, i/\mathcal{F}^{t-1}) = \underbrace{p(\{X\}/i, \mathcal{F}^{t-1})}_{\text{MT-pdf underlying}H_i} \underbrace{p(i/\mathcal{F}^{t-1})}_{w_i=\text{prob. of}H_i}, \quad (1)$$

where MT-pdf is the multi-object pdf underlying a hypothesis. Given the hypothesis, the MT-pdf is a product of independent individual pdfs underlying the objects, i.e.,

$$p(\{X\}/i, \mathcal{F}^{t-1}) = \prod_{k=1}^{M} p_k(x_k), \quad (2)$$

where $p_k(.)$ is the pdf of the k^{th} object.

Remark 1: In random finite set theory, the arguments of the MT-pdf $\{x_1, x_2 \cdots x_M\}$ above are interchangeable and thus, the MT-pdf is represented as:

$$p(\{X\}/i, \mathcal{F}^{t-1}) = \sum_{\bar{\sigma}} \prod_{k=1}^{M} p_{\sigma_k}(x_k), \quad (3)$$

where $\bar{\sigma} = \{\sigma_1, \sigma_2 \ldots \sigma_M\}$ represents all possible permutations of the indices $\{1, 2 \cdots M\}$. Hence, in any integration involving such a set of indices, a normalizing factor of $\frac{1}{M!}$ is used. In our case, we explicitly assign the index x_k to the target k or more precisely, the k^{th} component of the MT-pdf, $p_k(.)$. Note that such an assignment is always possible and there is no information loss in such a representation. Moreover, at the expense of more bookkeeping, this allows us to keep track of the labels of the different components of our multi-target pdfs. Please see the reference [19] where we show the equivalence of the hypothesis level equations derived here and the FISST recursions.

Next, we consider the prediction step between measurements. Each M-object hypothesis H_i splits into A_M children hypotheses

$$A_M = \sum_{n=0}^{min(m,M)} \binom{M}{n} \binom{m}{n} n!, \quad (4)$$

where m is the number of measurement returns. We note here that this is a pseudo-prediction step since we assume that we know the size of the return m. However, it allows us to fit the MT-tracking method nicely into a typical filtering framework. Using the rules of total and conditional probability, it follows that the predicted multi-object pdf in terms of the children hypotheses is:

$$p^-(\{X\}, (i,j)/\mathcal{F}^{t-1}) = \underbrace{\int p(\{X\}/(i,j), \{X'\}) p(\{X'\}/i, \mathcal{F}^{t-1}) d\{X'\}}_{p^-(\{X\}/(i,j), \mathcal{F}^{t-1})} \underbrace{p(j/i)}_{p_{ij}} \underbrace{p(i/\mathcal{F}^{t-1})}_{w_i}, \quad (5)$$

where $p^-(.,(i,j)/\mathcal{F}^{t-1})$ is the joint distribution of the state and hypothesis pairs before the measurement at time t, p_{ij} is the transition probability of going from the parent i to the child j and w_i is the probability of the parent hypothesis H_i. Let $p_k(x_k/x'_k)$ denote the transition density function of the k^{th} object.

$$p^-(\{X\}/(i,j),\mathcal{F}^{t-1}) \equiv \\ = \prod_k \int p_k(x_k/x'_k)p_k(x'_k)dx'_k = \prod_k p_k^-(x_k), \quad (6)$$

where $p_k^-(x_k)$ is the prediction of the k^{th} object pdf underlying the hypothesis H_{ij}. If a priori information, for instance, in terms of a probability of detection p_D is available, then:

$$p_{ij} = \frac{p_D^k(1-p_D)^{M-k}}{\binom{m}{k}k!}, \quad (7)$$

where ij is a data association in which exactly k of the M targets are associated to measurements. The $\binom{m}{k}k!$ factor is required so that p_{ij} is a valid probability distribution , i.e., $\sum_j p_{ij} = 1$.

Given the prediction step above, let us consider the update step given the measurements $\{Z_t\} = \{z_{1,t}, \cdots z_{m,t}\}$. Using Bayes rule:

$$p(\{X\},(i,j)/\mathcal{F}^t) = \\ \frac{p(\{Z_t\}/\{X\},(i,j))p^-(\{X\}/(i,j),\mathcal{F}^{t-1})}{l_{ij}} \frac{l_{ij}\overbrace{p_{ij}w_i}^{w_{ij}}}{\sum_{i',j'} l_{i',j'}\underbrace{p_{i'j'}w_{i'}}_{w_{i'j'}}}, \quad (8)$$

where

$$l_{ij} = \int p(\{Z_t\}/\{X'\},(i,j))p^-(\{X'\}/(i,j),\mathcal{F}^{t-1})d\{X'\}. \quad (9)$$

Note that l_{ij} is likelihood of the data $\{Z_t\}$ given the multi-object pdf underlying hypothesis H_{ij}, and the particular data association that is encoded in the hypothesis, i.e., $z_i \mapsto x_{ji}$. Then,

$$l_{ij} = \prod_{i=1}^m p(z_i/X_{j_i}), \quad (10)$$

where

$$p(z_i/X_{j_i}) = \begin{cases} \int p(z_i/x)p_{j_i}(x)dx & \text{if } X_{j_i} \in \mathcal{T} \\ g(z_i) & \text{if } X_{j_i} \in \mathcal{C} \end{cases} \quad (11)$$

where $\mathcal{T}$ is the set of all objects and $\mathcal{C}$ is clutter and the above equation implies that the measurement z_i was associated to clutter if $X_{j_i} \in \mathcal{C}$. Thus, given that the likelihoods of different hypothesis l_{ij} arise from the underlying multi-object pdf and the encoded data association in the hypotheses H_{ij}, the FISST updates can be written purely at the hypothesis level as follows:

$$w_{ij} := \frac{l_{ij}w_{ij}}{\sum_{i',j'} l_{i'j'}w_{i'j'}}, \quad (12)$$

where $w_{ij} = p_{ij}w_i$. Thus, we can see that the FISST update has a particularly simple Bayesian recursive form when viewed at the discrete hypothesis level, given that the multi-object pdfs underlying the hypotheses H_{ij} are tracked using some suitable method. The hypotheses H_{ij} now become the parent hypotheses for the next time step.

B. Relationship to MHT

The MHT likelihood for a child hypothesis j of parent i is of the form:

$$\eta_{ij}^{MHT}(z_1,..z_m) = p_D^k(1-p_D)^{M-k}\prod_{l=1}^m p(z_l/\hat{x}_{j_l}), \quad (13)$$

where $\hat{x}_{j_l}$ is the mean of the pdf of the target X_{j_l} or $p(z_k/\hat{x}_{j_l}) = g(z_l)$ if measurement z_l is associated to clutter. The equivalent hypothesis likelihood in the hypothesis level FISST (H-FISST) derivation is:

$$\eta_{ij}^{HFISST}(z_1,..z_m) = \\ p_{ij}l_{ij} = \frac{{p_D}^k(1-p_D)^{M-k}}{\binom{m}{k}k!}\prod_{l=1}^m p(z_l/X_{j_l}), \quad (14)$$

where the terms under the product in the above equation have been defined in Eq. 10. Thus, it may be seen that the main difference in the two likelihoods is the factor $\binom{m}{k}k!$ and the fact that $p(z_l/\hat{x}_{j_l})$ is an approximation of $p(z_l/X_{j_l})$ for observations that have been associated to a target. The factor is required such that the likelihood is normalized. The normalization of the likelihood, as in HFISST, is necessary from a probabilistic perspective since otherwise the distribution on the filtered state pdf, resulting from all possible observations, does not constitute a probability distribution, i.e., it does not add up to unity. This can easily be seen for the case of a standard filtering problem which carries over to the mutli-target tracking problem. Let the filtered state pdf, the belief state, be denoted by $b(x)$. Suppose that the likelihood function $\int p(z/x)dz \neq 1$. Consider the distribution on the future belief state $b'(x)$. This is given by:

$$p(b'/b) = \int_z p(b'/z,b)p(z/b)dz, \text{where} \\ p(z/b) = \int p(z/x)b(x)dx. \quad (15)$$

Note that if $\int p(z/x)dz \neq 1$ then $\int p(z/b)dz \neq 1$. Hence,

$$\int p(b'/b)db' = \int\int p(b'/z,b)p(z/b)dzdb' \\ = \int p(z/b)dz \neq 1. \quad (16)$$

We know that the filtered pdf (the belief process) has to evolve according to a Markov chain [20], [21] but the above development shows that the evolution equation violates the

requirement that the transition probability of a Markov chain needs to be a probability distribution over all future states, if the likelihood does not normalize to unity.

C. Equivalence of MHT and FISST

The proposed hybrid derivation (in that it includes both the continuous and discrete parts of the problem) as opposed to MHT which is a purely discrete derivation at the hypothesis level [13], reveals the critical hybrid structure (Eq. 8) inherent to multi-target tracking problems, and which, in turn allows us to unify the HFISST development with the FISST based formulation of the multi-target tracking problem, and thus, allows for the unification of FISST and MHT: methods that have thus far been thought to be different from each other. Because of the paucity of space, we cannot reproduce the derivation here, but the reader is referred to the technical report [19] for more details.

III. A Randomized FISST (R-FISST) Technique

In this section, we show how birth and death processes can be conveniently included in the HFISST technique, and propose a randomized implementation of the same, termed Randomized FISST (RFISST).

A. Incorporating Birth and Death Processes into HFISST

We assume the following model for the brith and death of targets.

a) Birth and Death Process:: We assume that the births and deaths can only happen in the field of view (FOV) of the sensor. This is done to ensure computational tractability but can be relaxed for theoretical purposes. We further assume that the FOV has been discretized into N pixels. Further, let us consider an M-target hypothesis H_i. We assume that the birth process is a simple Binomial process where the probability of a birth in a given pixel at a given time instant is α, and this is independent of the birth in any other pixel. Further, we assume that the death process is also Binomial, and any target in the FOV can disappear with a probability β independent of the other targets in the FOV. Moreover, the birth and death processes are independent of each other.

If we let the number of pixels tend to infinity, and we let the number of time instants, n, go to infinity, such that $n\alpha = \lambda$, we recover a Poisson Point process.

It may be seen that the probability of any instance of exactly N_b births and N_d deaths starting with the M-target hypothesis given by:

$$p_{ij} = \alpha^{N_b}\beta^{N_d}, \tag{17}$$

and the number of such instances is $\binom{N}{N_b}\binom{M}{N_d}$. Given the child H_{ij} of the parent H_i that incorporates a particular number of birth and death hypothesis, due to the data association, the child H_{ij} can further split into grandchildren H_{ijk} where, from before, the probability of the k^{th} data association hypothesis given the ij^{th} child is:

$$p_{ij,k} = \frac{p_D^l(1-p_D)^{m-l}}{\binom{m}{l}l!}, \tag{18}$$

where m is the number of returns, and the ijk^{th} child corresponds to a data association hypothesis that chooses to associate l of the returns to objects in the FOV. Thus, given a particular M object hypothesis, the total number of possible hypotheses (at the grandchild level) after receiving a measurement is:

$$\widetilde{A}_M = \sum_{K=0}^{N} a_{M+K}A_{M+K} + \sum_{K=-1}^{M} a_{M-K}A_{M-K}, \tag{19}$$

where A_{M+K} is calculated using Eq. 4 and,

$$a_{M+K} = \sum_{j=0}^{N}\binom{N}{K+j}\binom{M}{j}, \tag{20}$$

$$a_{M-K} = \sum_{j=0}^{M-K}\binom{M}{K+j}\binom{N}{j}. \tag{21}$$

The birth hypothesis includes a continuous pdf for any birthed object to form the MT-pdf underlying the birth hypothesis, we do not go into the details here but this may be done in a straightforward, albeit tedious, fashion and is shown in our previous work [1], [2]. Thus, the probability of the jk^{th} child given the parent hypothesis H_i is simply:

$$p_{i,jk} = p_{ij}p_{ij,k}. \tag{22}$$

The likelihoods of any of these hypothesis, l_{ijk}, can now be calculated using Eq. 10 from before. Noting that jk can be replaced by a single number l, it follows that the HFISST procedure carries through to the birth/ death case with minimal changes, except now the transition probability for a particular (grand)child hypothesis is given by Eq. 22. However, note that the number of (grand)children hypothesis explodes further due to the birth and death process. Please refer to Fig. 1 for an illustration of the process.

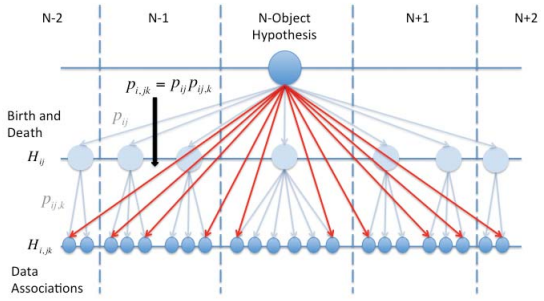

Fig. 1. An illustration of the multiple birth and death process. Arrows show connection between parent, child, and grandchild hypotheses.

B. MCMC based Intelligent Sampling of Children Hypothesis

Recall Eq. 12. It is practically plausible that most children j of hypothesis H_i are highly unlikely, i.e., $l_{ij} \approx 0$ and thus, $w_{ij} \approx 0$. Hence, there is a need to sample the children H_{ij} of hypothesis H_i such that only the highly likely hypotheses are sampled, i.e., $l_{ij} >> 0$.

Remark 2: Searching through the space of all possibly hypotheses quickly becomes intractable as the number of objects and measurements increase, and int he presence of multiple birth and deaths. We cannot sample the hypothesis naively either, for instance, according to a uniform distribution, since the highly likely hypothesis are very rare under the uniform distribution, and thus, our probability of sampling a likely hypothesis is vanishingly small under a uniform sampling distribution. This is the equivalent to the "particle depletion" problem seen in particle filters.

Thus, we have to resort to an intelligent sampling technique. In particular, given a hypothesis H_i, we want to sample its children according to the probabilities $\bar{p}_{ij} = w_{ij} l_{ij}$. This can be done by generating an MCMC simulation where the sampling Markov chain, after enough time has passed (the burn in period), will sample the children hypotheses according to the probabilities $\bar{p}_{ij}$. A pseudo-code for setting up such an MCMC simulation is shown in Algorithm 1. In the limit,

Algorithm 1 MCMC Hypothesis Sampling

Generate child hypothesis j_0, set $k = 0$.
Generate $j_{k+1} = \pi(j_k)$ where $\pi(.)$ is a symmetric proposal distribution
If $\bar{p}_{ij_{k+1}} > \bar{p}_{ij_k}$ then $j_k := j_{k+1}; k := k + 1;$
else $j_k := j_{k+1}$ with probability proportional to $\frac{\bar{p}_{ij_{k+1}}}{\bar{p}_{ij_k}}$; $k = k + 1$.

as $k \rightarrow \infty$, the sequence $\{j_k\}$ generated by the MCMC procedure above would sample the children hypotheses according to the probabilities $\bar{p}_{ij}$. Let us keep the highest likely C children hypothesis for every hypothesis. Given these $M*C$ hypotheses, i.e. C children of M parents, we can keep a fixed number H_∞ at every generation by either sampling the H_∞ highest weighted hypotheses among the children, or randomly sampling H_∞ hypotheses from all the children hypotheses according to the probabilities w_{ij}.

Remark 3: It may be shown that the mixing time of an MCMC chain is $\mathcal{O}(logN)$ where N is the number of states in the chain [22] if the Markov Chain's "congestion" is suitably bounded. The MCMC simulation below has this property, and thus, can scale to situations with a very large number of hypotheses. Due to the paucity of space, we do not show this here but this will be the subject of a forthcoming expanded journal version of the current paper.

Now, we show the details of how to implement the multiple birth and death randomized hypothesis generation technique. Like our previously published papers, [1], [2], this method takes advantage of MCMC at every time step to generate only the highly probable hypotheses. This method differs from previous methods because it accounts for multiple birth and death scenarios at each time step. Consider the scenario shown in figure 2. In this scenario, a single hypothesis predicts there are ten objects in the field of view, which are represented by circles. A single sensor with a thirty-degree field of view takes a measurement and receives five returns represented by diamonds. Knowing the number of objects $M = 10$ and the number of measurements $m = 5$ the total number of possible hypotheses can be calculated using Eq. (4), $A_M = 63,591$. If we were to exhaustively generate the possible hypotheses

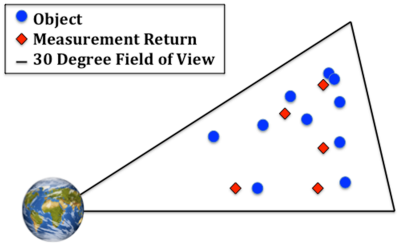

Fig. 2. A possible SSA event where there is assumed to be ten objects in the field of view and five measurement returns.

from this particular scenario we would start by constructing a matrix called the hypothesis matrix. Each row of this matrix would represent a unique hypothesis. Generating this hypothesis matrix would become increasingly difficult as the number of objects and measurement returns increased. Even for a simple sixty object example with one spawn occurrence the maximum number of possible hypotheses at one time step can be in the order of tens of billions. Our randomized approach allows us to avoid generating this matrix in its entirety. Instead we only generate the rows of the hypothesis matrix that correspond to highly probable hypotheses. By taking each row of the hypothesis matrix as the state of a Markov Chain and we explore the chain using an MCMC criterion that favors the highly probable states. We can do this tractably by generating a matrix called the data association matrix (figure 3). Creating the data association matrix has no added computational cost because each element of the matrix is necessary in calculating hypothesis probability and is used in every method of hypothesis generation. Also, the dimensions of the data association matrix are much smaller than those of the hypothesis matrix. This makes it much more practical to explore using MCMC. The data association matrix for the new randomized multiple birth and death approach has dimensions $(m + 1) \times (M + 1 + 1)$. That being said each row of the matrix represents a measurement return plus

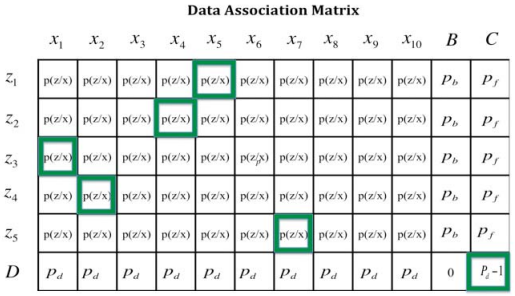

Fig. 3. The Data Association Matrix. Each row represents a particular measurement return. Each column is a particular association. The elements represent the likelihood of the corresponding association. Green boxes here show a visual representation of an example hypothesis.

an added row to represent death. The columns represent the possible associations including a column for each associable object, a column for association to birth (B), and a column for no association (C). Each element of the matrix is the corresponding measurement to association likelihood.

Remark 4: The numbering of the objects and measurement returns in the data association matrix is done strictly for organization and is redone at random each time step with no record of previous numbering or labeling kept throughout scans.

We start the MCMC procedure with a single row vector of length $m + 1$ containing a random permutation of the numbers 1 through $(M + 2)$ and the letter N to signify a no death hypothesis. For example consider such a row vector $[5\ 4\ 1\ 2\ 7\ N]$.This row vector is our current hypothesis and is represented on the data association matrix in figure 3 using boxes. We then propose a new hypothesis to walk to in the MCMC. We do this by randomly choosing a measurement return and switch its association using a uniform proposal distribution. We make sure the proposed switch creates a valid hypothesis with no conflicting associations. If there is a conflicting association we automatically assign the conflict to clutter. The resulting hypothesis becomes our proposed hypothesis. These steps are displayed in figure 4.

Remark 5: Conflicting associations can be seen as more than one box in a single column. Columns B and C are an exception to this rule because there can be multiple associations to birth as well as clutter in a single hypothesis. Multiple birth hypotheses are developed naturally in the MCMC when more than one measurement to birth association probability is higher than or comparitive to the corresponding object associations. We compare the proposed hypothesis to the current hypothesis using an MCMC criterion that stems from the Metropolis condition $U[0,1] < min(1, \frac{\bar{p}_{ij_{k+1}}}{\bar{p}_{ij_k}})$ where $\bar{p}_{ij_k}$ is the probability

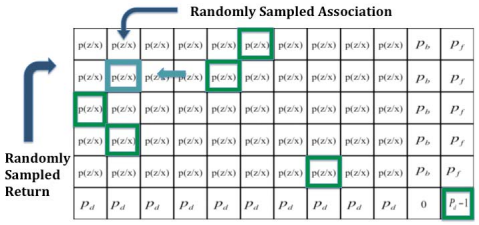

(a) Choose a random return and association to switch

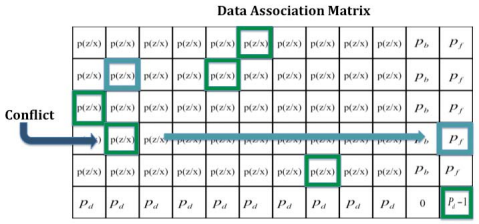

(b) Check for conflicting assignment. Assign conflict to clutter if necessary

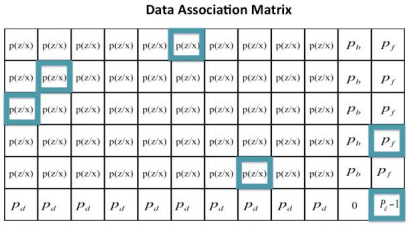

(c) New proposed hypothesis

Fig. 4. Visualization of a single MCMC step using the Data Association Matrix. This particular example contains a conflicting assignment with measurement return two and shows how the association is then changed to clutter.

of the hypothesis at step k. Simply, if the proposed hypothesis has a higher probability then it becomes the current hypothesis, if not, it becomes the current hypothesis with probability proportional to the ratio of the hypothesis probabilities. These steps are then repeated until we have reached an assumed stationary distribution. We then continue walking for a user defined amount of steps and record all the repeated hypotheses.

The recorded hypotheses are those with very high probability.

IV. Applications

In this section we apply the RFISST technique to a multi-object Space Situational Awareness (SSA) tracking problem. This simulation includes the birth of objects in the form of spawning to show how the method performs when there are multiple ambiguous measurement returns. Our goal is to highlight areas where the R-FISST approach is not only novel but neccessary to achieve accurate tracking data. We provide results from a twenty object Space Situational Awareness problem with a single spawn occurrence. We also show the growth of the number of possible hypotheses using a larger sixty object SSA example.

A. R-Fisst Application to Multi-Object Tracking Scenario

We will test the algorithm developed in this paper by applying it to the SSA multi-object tracking problem. To show an instance were the randomized approach will be most useful we developed scenarios in which spawning occurs. We consider a spawn to be an event in which a known space object breaks apart into multiple objects. An example of such event is shown if figure 5.

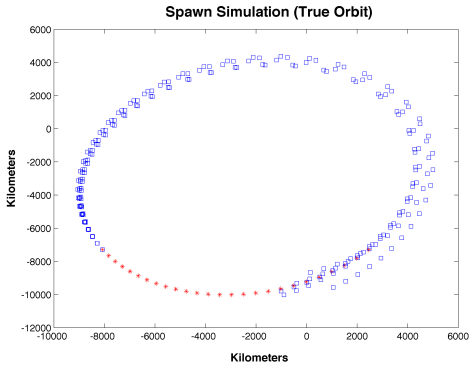

Fig. 5. Spawn simulation. A known object marked as red stars randomly splits into multiple objects (shown here as blue squares)

These spawn events are of particular interest to our research because they create a scenario with multiple ambiguous measurement returns and lead to large numbers of birth and death hypotheses. This causes a rise in the total number of possible hypotheses. As a proof of concept we created a simple SSA tracking problem of a single object depicted in figure 6. The initial state of this object was known with some uncertainty. A single sensor with a field of view of thirty degrees was used to capture measurements of the objects x and y position. At a random time in the orbit the object undergoes a spawn event. The goal of the simulation was to accurately estimate the states of all the newly spawned objects. Figure 6 shows the actual states of the spawned objects as they pass through the field of view next to the estimated positions from the top hypothesis. It can be seen from the figures that the R-FISST methodology was able to accurately predict the birth and track all the spawned objects. It is important to note that the method also correctly predicted that the original object had died. If it did not predict a death, the hypothesis would contain the wrong number of objects.

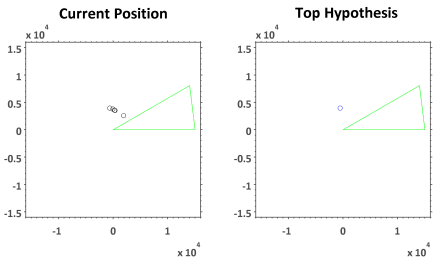

(a) Before measurements.

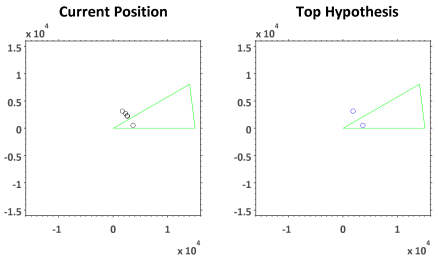

(b) Partially in the field of view.

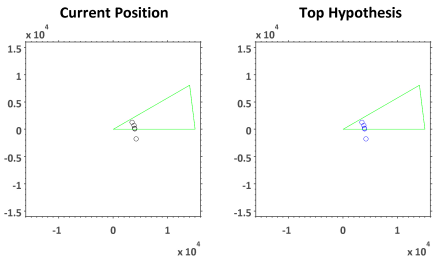

(c) Completely in the field of view.

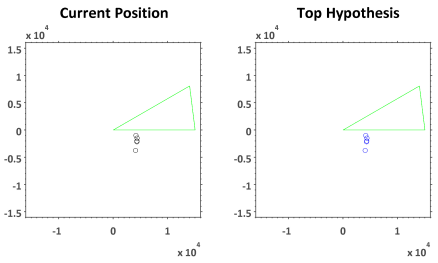

(d) After passing through the field of view.

Fig. 6. Snapshots of the actual states (black) and the estimated states from the top hypotheses (green) while passing through the field of view. Axes in tens of thousands of kilometers

We ran a similar simulation to determine if the methodology could capture a birth amongst other non-spawning objects. We did this by simulating a twenty-object SSA example. The initial states for all objects are known with some uncertainty and the same sensor was used. During the simulation one of the twenty objects spawns into multiple objects (shown in red). The goal of this simulation was to accurately track all new object while maintaining hypotheses with the correct number of objects. This simulation is shown in figure 7. From figure 7(a) one can see that before the spawned object enters the field of view it is still assumed to be intact. As the spawned objects enter the field of view the top hypothesis begins to predict the birth of the spawned objects until all objects are captured. The user defined α and β values must be tuned to accurately capture this scenario. In our simulation, heuristics were embedded in the code to adjust the α and β values dependent on the ratio between the number of measurements and the number of objects in the field of view. A rigorous approach for adjusting the α and β values online is a current topic of our research and will be addressed in future work.

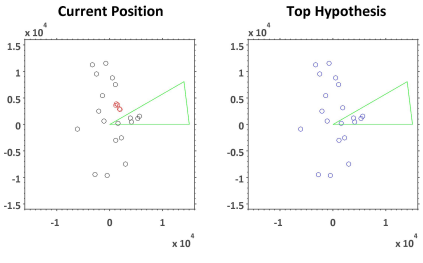

(a) Before measurements.

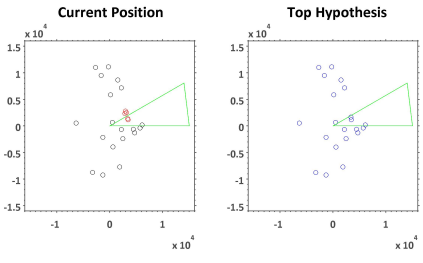

(b) Partially in the field of view.

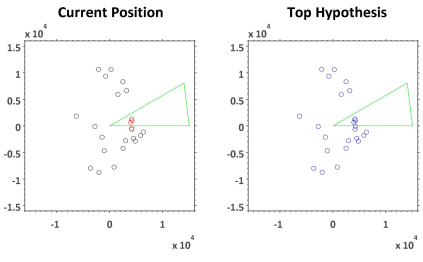

(c) Completely in the field of view.

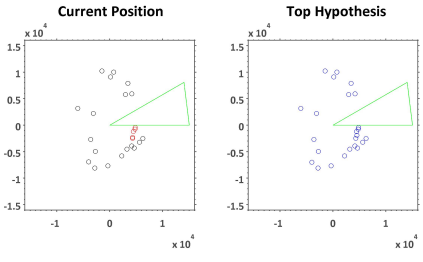

(d) After passing through the field of view.

Fig. 7. Snapshots of the actual states (black and red) and the estimated states from the top hypotheses (green) while passing through the field of view. Axes in tens of thousands of kilometers

To show the effect of the combination of the spawn event

and the ambiguous measurement returns, figure 8 displays the number of possible hypotheses throughout the simulation.

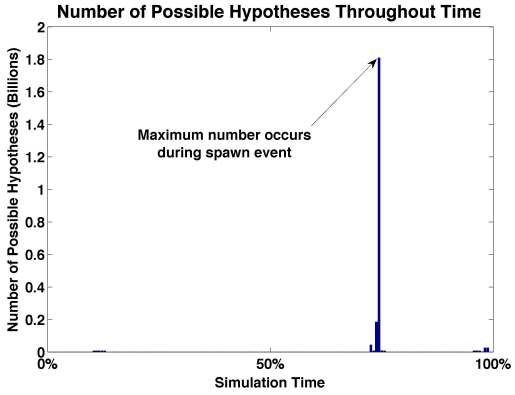

Fig. 8. Maximum number of possible hypotheses at each step throughout the entire simulation.

Note that the maximum number of possible hypotheses occurs when the spawn event passes through the field of view. This is because the number of possible hypotheses is a function of the number of associable objects and the number of measurement returns. The more ambiguous measurements the greater the number of hypotheses which in this moderate sized case still runs into the billions. Thus, exhaustively generating these many hypotheses is computationally intractable using conventional methods and emphasizes the need for a computationally tractable approach like RFISST.

V. CONCLUSIONS

In this paper, we have presented a newly developed randomized approach to the multi-object tracking problem called RFISST. RFISST provides a tractable solution to situations were ambiguous measurement returns cause the number of possible hypotheses to increase and become computationally burdensome. We applied the RFISST technique to SSA tracking problems with a random spawn event. The RFISST technique was able to track the spawn event and predict the correct number of objects even though the possible hypotheses were, at times, too numerous to generate exhaustively. In future work we will scale the problem to larger SSA applications, with multiple spawn events. We will also develop a rigorous approach for adjusting the α and β values online, which was done heuristically in this paper.

VI. ACKNOWLEDGMENTS

This work is funded by AFOSR grant number: FA9550-13-1-0074 under the Dynamic Data Driven Application Systems (DDDAS) program.

REFERENCES

[1] W. Faber, S. Chakravorty, and I. Hussein, "A randomized sampling based approach to multi-object tracking," *International Conference of Information Fusion,*, July 2015.

[2] ——, "A randomized sampling based approach to multi-object tracking with comparison to homht," *IAAS/AIAA Astrodynamics Specialists Conference,*, August 2015.

[3] I. R. Goodman, R. P. S. Mahler, and H. T. Nguyen, *Mathematics of Data Fusion*. Kluwer Academic Publishers, 1997.

[4] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artec House, 2007.

[5] B. Vo and W. Ma, "The gaussian mixture probability hypothesis density filter," *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, vol. 54, no. 11, November 2006.

[6] B. Vo *et al.*, "Sequential monte carlo methods for multi target filtering with random finite sets," *IEEE Tr. Aerosp. Electronic Systems*, vol. 4, pp. 1224–1245, 2005.

[7] B. T. Vo *et al.*, "An analytic implementation of the cardinalized probability hypothesis density filter," *IEEE Tr. Signal Processing*, vol. 55, pp. 3553–3567, 2007.

[8] B. Ristic *et al.*, "Improved smc implementation of the phd filter," in *PRoc. Conf. Info. fusion (FUSION)*, 2010.

[9] I. Hussein, K. DeMars, C. Frueh, R. S. Erwin, and M. Jah, "An aegis-fisst algorithm for joint detection and tracking in space situational awareness," *AAS/AIAA Astrodynamics Specialist Conference*, August 2012.

[10] ——, "An aegis-fisst integrated detection and tracking approach to space situational awareness," *International Conference of Information Fusion*, July 2012.

[11] B.-N. Vo *et al.*, "Labeled random finite sets and multi-object conjugate priors," *IEEE Transactions on Signal Processing*, vol. 61, pp. 3460–3475, 2013.

[12] ——, "Labeled random finite sets and the bayes multi-target tracking filter," *IEEE Transactions on Signal Processing*, vol. 62, pp. 6554–6567, 2014.

[13] D. B. Reid, "An algorithm for tracking multiple targets," *IEEE Transactions on Automatic Control*, vol. 24, pp. 843–854, 1979.

[14] Y. B. Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*, New York, 1995.

[15] Y. Blair and E. W. D. Blair, *Multitarget Multisensor tracking: Applications and advances, vol. III*. Norwood, MA: Artech House, 2000.

[16] S. Oh *et al.*, "Markov chain monte carlo data association for multi-target tracking," *IEEE Tr. Automatic Control*, vol. 54, pp. 481–497, 2009.

[17] S. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA: Artech House, 1999.

[18] Y. B. Shalom and T. Fortmann, *Tracking and Data Association*. San Diego, CA: Academic Press, 1988.

[19] W. Faber, S. Chakravorty, and I. I. Hussein, "Mutli-target tracking with a randomized hypothesis generation technique," *IEEE transaction on Aerospace And Electronic Systems*, 2016, under review, preprint available at ArXiv identifier 1603.04096.

[20] D. P. Bertsekas, *Dynamic Programming and Optimal Control, vols I and II*. Cambridge: Athena Scientific, 2000.

[21] P. R. Kumar and P. P. Varaiya, *Stochastic Systems: Estimation, Identification and Adaptive Control*. Prentic Hall, NJ: Prentice Hall, 1986.

[22] M. Jerrum and A. Sinclair, "The markov chain monte carlo method: An approach to approximate counting and integration," in *Approximation Algorithms for NP hard Problems, D. S. Hochbaum Ed.*, 1996.